



(11)

EP 1 202 597 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
02.05.2002 Bulletin 2002/18

(51) Int Cl.7: **H04Q 11/00**, **H04L 12/56**

(21) Application number: 01309176.4

(22) Date of filing: 30.10.2001

(84) Designated Contracting States:
**AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
 MC NL PT SE TR**
 Designated Extension States:
AL LT LV MK RO SI

(30) Priority: 31.10.2000 US 703064

(71) Applicant: **Chiaro Networks Ltd.**
Richardson, Texas 75080 (US)

(72) Inventors:

- **Brewer, Tony M.**
Plano, Texas 75093 (US)
- **Blackmon, Harry C.**
Plano, Texas 75023 (US)
- **Dozier, Harold W.**
Dallas, Texas 75248 (US)

- Kleiner, Jim
Dallas, Texas 75240 (US)
- McDermott, Thomas C., III
Plano, Texas 75094 (US)
- Palmer, Gregory S.
Plano, Texas 75025 (US)
- Shaw, Keith W.
Plano, Texas 75023 (US)
- Traylor, David
Richardson, Texas 75080 (US)
- Walker, Dean E.
Plano, Texas 75025 (US)

(74) Representative: Jackson, Richard Eric et al
Carpmaels & Ransford,
43 Bloomsbury Square
London WC1A 2RA (GB)

(54) Router-Switch fabric protection using forward error correction

(57) Instead of alternatively utilizing only one fabric or the other fabric of a redundant pair, both fabrics simultaneously transmit duplicate information, such that each packet forwarding module (PFM) receives the output of both fabrics simultaneously. In real time, an internal optics module (IOM) analyzes each information chunk coming out of a working zero switch fabric; simultaneously examines a parallel output of a working one duplicate switch fabric; and compares on a chunk-by-chunk basis the validity of each and every chunk from both switch fabrics. The IOM does this by examining forward error correction (FEC) check symbols encapsulated into each chunk. FEC check symbols allow correcting a predetermined number of bit errors within a chunk. If the chunk cannot be corrected, then the IOM provides indication to all PFMs downstream that the chunk is defective. Under such conditions, the PFMs select a chunk from the non-defective switch fabric. Under error-free normal conditions, however, the PFMs select a chunk arbitrarily from a default switch fabric. In this way, each chunk in real time is selected from a non-defective source and is thus guaranteed to be error free. Accordingly, if a switch fabric fails, no information chunks are lost anywhere in the system.






FIG. 1A



Description

RELATED APPLICATIONS

[0001] This application is related to concurrently filed, co-pending, and commonly assigned U.S. Application Serial Number 09/703,057, entitled "System And Method For IP Router With an Optical Core," to concurrently filed, co-pending, and commonly assigned U.S. Application Serial Number 09/703,056, entitled "System and Method for Router Central Arbitration," to concurrently filed, co-pending, and commonly assigned U.S. Application Serial Number 09/703,038, entitled "System and Method for Router Data Aggregation and Delivery," to concurrently filed, co-pending, and commonly assigned U.S. Application Serial Number 09/702,958, entitled "Timing and Synchronization for an IP Router Using an Optical Switch," to concurrently filed, co-pending, and commonly assigned U.S. Application Serial Number 09/703,027, entitled "Router Network Protection Using Multiple Facility Interfaces," and to concurrently filed, co-pending, and commonly assigned U.S. Application Serial Number 09/703,043, entitled "Router Line Card Protection Using One-for-N Redundancy," the disclosures of which are incorporated herein by reference.

TECHNICAL FIELD

[0002] This application relates to the field of optical communication networks, and particularly to large-scale routers for optical communication networks.

BACKGROUND

[0003] Routers form a central part of a data communication network and perform general routing of data packets. Headers in each packet control the route taken by that packet through the network. There can be multiple routers in a network. Information, as contained in the data packets, typically travels from one router to the next router, and eventually reaches the destination edge of the network, where a destination edge router receives the information packet and decides where it goes from there. Typically it goes to an Internet service provider at the opposite edge of the edge router. If the destination is a household PC, the Internet service provider then sends the information to the destination computer. If there is corporate access to the network, the information may go from the edge router directly to a corporate site.

[0004] A fabric is a collection of devices which cooperatively provide a general routing capability. Internet protocol (IP) routers require protection from fabric failures, for example optical fabric, packet fabric, and switch element fabric failures. The prior art uses duplicated switch fabrics and line cards that feed both switch fabrics simultaneously but receive from only one switch fabric at any given time.

[0005] Prior art designs based upon duplex or dupli-

cated fabrics, while fully protecting from individual fabric failures, permit extensive loss of packets during a failure, detection, and protection event. In the prior art a line card receives output from one switch fabric, unless it detects a failure on that switch fabric, in which case it then spends about 50 milliseconds switching over to the redundant switch fabric. In an outage of a fabric, if it requires fifty milliseconds for the receiving line card to determine that a fabric is defective and to decide to start receiving instead from the redundant fabric, then all packets passing through that defective fabric during that fifty milliseconds are lost. The number of packets lost per line card is equal to the switchover time to the alternate fabric multiplied by the packet transmission rate through that particular switch fabric. Multiplying by the total number of line cards in the system, the result can be a huge amount of data lost during each 50 millisecond transient event, which creates a large impact on the network as a whole.

[0006] Needed in the art are a new system and method of IP router switch fabric protection that prevent or minimize loss of data packets during a fabric failure, detection, and protection event.

SUMMARY OF THE INVENTION

[0007] The present invention is directed to a system and method which, instead of selectively utilizing one fabric or the other fabric of a redundant pair, utilize both fabrics simultaneously and transmit duplicate identical information through both fabrics, such that each packet forwarding module (PFM) receives the output of both fabrics simultaneously.

[0008] All information moves through an optical switch fabric in chunks. A chunk is defined as a uniformly sized unit of information that is passed through an optical switch during one cycle of the switch. In real time, an internal optics module (IOM) analyzes each chunk coming out of a working zero switch fabric; simultaneously examines the output of a working one copy of this switch fabric; and compares on a chunk-by-chunk basis the validity of each and every chunk from both switch fabrics. The IOM does this by examining forward error correction (FEC) check symbols encapsulated into each chunk. The FEC check symbols allow correcting a predetermined number of bit errors within a chunk. If the chunk cannot be corrected, then the IOM provides indication to all PFMs downstream from the IOM that the chunk is defective. Under such conditions, the PFMs select a chunk from the non-defective switch fabric. Under error-free normal conditions, however, the PFMs select a chunk arbitrarily from a default switch fabric.

[0009] Thus, a PFM receives chunks from both working zero and working one copies of a switch fabric. If both chunks are error-free, it arbitrarily discards one chunk and keeps the duplicate chunk, since either chunk is good. If the PFM receives two simultaneous duplicate chunks, and one of those chunks is defective,

the PFM recognizes and discards the defective chunk and keeps the good chunk. If both simultaneous chunks are defective, then a double fault has occurred, which is an unlikely occurrence and is outside the principal objects of the present invention.

[0010] In this way, each chunk in real time is selected from a non-defective source and is thus guaranteed to be error free. Accordingly, if a switch fabric fails, no chunks are lost anywhere in the system.

[0011] Various aspects of the invention are described in concurrently filed, co-pending, and commonly assigned U.S. Application Serial Number 09/703,057, entitled "System And Method For IP Router With an Optical Core," concurrently filed, co-pending, and commonly assigned U.S. Application Serial Number 09/703,056, entitled "System and Method for Router Central Arbitration," concurrently filed, co-pending, and commonly assigned U.S. Application Serial Number 09/703,038, entitled "System and Method for Router Data Aggregation and Delivery," concurrently filed, co-pending, and commonly assigned U.S. Application Serial Number 09/702,958, entitled "Timing and Synchronization for an IP Router Using an Optical Switch," concurrently filed, co-pending, and commonly assigned U.S. Application Serial Number 09/703,027, entitled "Router Network Protection Using Multiple Facility Interfaces," and concurrently filed, co-pending, and commonly assigned U.S. Application Serial Number 09/703,043, entitled "Router Line Card Protection Using One-for-N Redundancy," the disclosures of which are incorporated herein by reference.

[0012] Embodiments according to the present invention are designed to protect against all single fault occurrences. Single faults include a single fault of a module, a single fault of a cable, or a single fault of a path. Accordingly, although some double faults are protected against, double faults generally lie beyond the scope of primary objects of the present invention and thus are not in general protected against.

[0013] The foregoing has outlined rather broadly the features and technical advantages of the present invention in order that the detailed description of the invention that follows may be better understood. Additional features and advantages of the invention will be described hereinafter which form the subject of the claims of the invention. It should be appreciated by those skilled in the art that the conception and specific embodiment disclosed may be readily utilized as a basis for modifying or designing other structures for carrying out the same purposes of the present invention. It should also be realized by those skilled in the art that such equivalent constructions do not depart from the spirit and scope of the invention as set forth in the appended claims. The novel features which are believed to be characteristic of the invention, both as to its organization and method of operation, together with further objects and advantages will be better understood from the following description when considered in connection with the accompanying

figures. It is to be expressly understood, however, that each of the figures is provided for the purpose of illustration and description only and is not intended as a definition of the limits of the present invention.

5

BRIEF DESCRIPTION OF THE DRAWING

[0014] For a more complete understanding of the present invention, reference is now made to the following descriptions taken in conjunction with the accompanying drawing, in which:

10 Figs. 1A-1C form a schematic diagram showing an overview of the data paths through a router, in an embodiment of the present invention;

15 Fig. 2 is a block diagram illustrating data flow through facility modules of a router in more detail;

20 Fig. 3 is a block diagram illustrating information flow through a typical packet forwarding module;

25 Fig. 4 is a block diagram representing information flow through a typical internal optics module, according to an embodiment of the present invention; and

30 Fig. 5 is a simplified schematic diagram illustrating a method of fault isolation and diagnostics in a switch fabric, in accordance with embodiments of the present invention.

DETAILED DESCRIPTION

35 [0015] Figs. 1A-1C form a schematic diagram showing an overview of the data paths through a router 10, in an embodiment of the present invention. For ease of understanding, Figs. 1A-1C are partitioned into three sequentially adjacent panels. Fig. 1A-1C do not show
40 how router system 10 is wired, but simply illustrates the flow of data. At the upper left portion of Fig. 1A, an input 101-0 is a first SONET data channel, formatted as Packet-over-SONET in the present embodiment. Input 101-0 includes two optical fibers, namely a working input fiber
45 101-0W and a protect input fiber 101-0P. Fibers 101-0W, 101-0P carry duplicated information into router 10 from a peer source equipment e.g., another router or piece of SONET transmission equipment, compatible with the Packet-over-SONET format. Protect and working facility
50 module cards 11-0P and 11-0W independently receive duplicate input from respective optic fibers 101-0P and 101-0W and perform an integrity check on the information by computing SONET parity and SONET framing words to determine if the information is valid, and inde-
55 pendently check SONET protection switching 'K' Bytes. Both facility modules 11-0W and 11-0P perform essentially identical functions on the information. Each facility module independently evaluates the SONET frame and

determines whether the information contained on it is valid. Facility modules 11-0W and 11-0P then extract packets from their respective SONET frames and transfer those packets over a packet bus 103 to a packet forwarding module (PFM) 13-0.

[0016] Working facility module 11-0W and protect facility module 11-0P respectively provide duplicate input interfaces 103-0W and 103-0P to packet forwarding module 13-0. A system controller (not shown in Figs. 1A-1C) examines the status of facility modules 11-0W and 11-0P and selects as the in-service facility module the one that is receiving properly framed and bit-error-free packets on its input, in accordance with the SONET standard or as selected by SONET 'K' Bytes. Should the information coming into a facility module, for example facility module 11-0P, have bit errors or other defects, then facility module 11-0P raises an alarm at the system controller (not shown in Figs. 1A-1C). The system controller then selects facility module 11-0W as the source of input from that channel, and facility module 11-0W strips the packets out of the SONET framing overhead and transfers those raw packets over industry standard bus 103-0W to packet forwarding module 13-0. Typically facility modules 11-0W and 11-0P, along with packet forwarding module 13-0, are contained in a line shelf, denoted in Fig. 1A as $\frac{1}{2}$ line shelf 142 (ingress) and described below in more detail.

[0017] There are actually N+1 multiple packet forwarding modules 13-0 through 13-N. In Fig. 1A, N equals 4, providing for four working packet forwarding modules 13-0 through 13-3 and a fifth designated protect packet forwarding module 13-4. In this case protect packet forwarding module 13-4 is a spare module available to replace any working module 13-0 through 13-3. Should any one of working packet forwarding modules 13-0 through 13-3 fail, then fifth packet forwarding module 13-4 can substitute for the failed packet forwarding module 13-0 through 13-3. This protection configuration is known as "one-for-four" protection. Similarly, on the output side of router 10 shown in the right side portion of Fig. 1C, packet forwarding modules 18-0 through 18-3 are all working modules, and packet forwarding module 18-4 is a spare protect packet forwarding module available as a replacement in the event of failure of any working packet forwarding module 18-0 through 18-3. Typically packet forwarding modules 18-0 through 18-4 are contained in a line shelf, denoted in Fig. 1C as $\frac{1}{2}$ line shelf 143 (egress) and described below in more detail.

[0018] Protection works through a daisy-chain data bus 105 cascading from Channel 0 to Channel 1, to Channel 2, to Channel 3, and to Channel 4, linking facility modules 11-0W through 11-4W. A duplicate data bus interconnects from Channel 4 up to Channel 0, linking facility modules 11-4P through 11-0P. If for example packet forwarding module 13-1 were to fail, then input facility modules 11-1P and 11-1W send their traffic down data bus 105 linking facility modules 11-2 and 11-3 to

facility module 11-4, which then switches those inputs to protect packet forwarding module 13-4. Thus if one channel fails, traffic, instead of going through the failed channel, goes down data bus chain 105 to designated protect module 13-4. If a particular facility module needs to be removed for maintenance purposes on one data bus, the duplicate data bus is maintained intact, allowing for hot replacement of any of the facility modules, working and protect, even if a packet forwarding module protection switch is in effect at the time. Similarly on the output side of router 10, output data is rerouted up a data bus chain 106 to Channel 1 and then out of router 10.

[0019] In operation, if PFM 13-1 fails, a microprocessor in the line shelf containing the failed packet forwarding module detects the failure, notices if the system is configured for one-for-four protection, and instructs switches on facility modules 11-1 through 11-4 to switch traffic that used to be in Channel 1 down to Channel 4. Channel 4 contains facility modules 11-4P and 11-4W on the input side and facility modules 12-4P and 12-4W on the output side respectively of router 10. These modules are connected to optical inputs and outputs only when utilizing protect PFM 13-4 or 18-4 as a working module and not as protection for PFMs 13-0 through 13-3 or 18-0 through 18-3. If PFM 13-4 or 18-4 is a working module, then daisy chain bus 105, 106 is not utilized in any way, and there are simply 5 working inputs and 5 working outputs. Accordingly, two modes of operation are available; namely one-for-N protection, for example one-for-four; or zero-for-five protection, meaning no protect modules and five working modules. Without requiring any wiring changes, router system 10 will function in either mode.

[0020] An alternative operating mode designates input 101-N and output 102-N for lower priority traffic. That traffic would be deliberately interrupted in the event of a failure of any of the packet forwarding modules carrying higher priority traffic and requiring a protect packet forwarding module to service that failure.

[0021] Information is transferred from PFM 13-0 to internal optics modules (IOMs) 14 as chunk payloads of data, such that a chunk contains typically 400 bytes of payload data. Packets contained in virtual out queues of PFM 13-0 that are destined for the same egress PFM can be combined to form a single chunk payload of data. Thus, multiple small packets or just a segment of a larger packet can be loaded into a single chunk. A maximum of two chunks can be transferred from a PFM 13-0 to the IOMs 14-0W0 and 14-1W0 during each chunk period. The same chunks are replicated and transferred in parallel to IOMs 14-0W1 and 14-1W1.

[0022] IOM modules 14 encapsulate FEC code words as multiple redundant check symbols into each of the chunks. The present implementation uses a conventional interleaved Reed-Solomon FEC coding. IO modules 14-0W0, 14-1W0 provide duplicate working module capacity for a working zero optical switch plane. Similarly

IO modules 14-0W1, 14-1W1 provide duplicate working module capacity for a working one optical switch plane. Switch plane pairs in this case are not configured as working and protect, but as working zero and working one copies respectively, such that copy zero switch plane containing optical switch modules 15-1 through 15-6 and duplicate copy one switch plane containing optical switch modules 16-1 through 16-6 each provide 6 optical switches worth of capacity.

[0023] IO module 14-0W0 transfers information from PFM 13-0 to one of three optical switch modules 15-1, 15-2 and 15-3. IO module 14-0W0 sends the information to the appropriate optical switch module based on the decisions of the central arbiter module (not shown in the figures), described in U.S. Application Serial Number 09/703,057, cited above. Illustratively, one input comes into an optical switch module and one output goes out from that same optical switch module. In an actual system, these inputs and outputs in fact provide connectivity across router system 10. Fig. 1B shows optical switch module 15-1 connected to an egress side internal optics module 17-0W0 through an output fiber 110-1. For clarity, six such optical switch modules 15-1 through 15-6 are shown in the top portion of Fig. 1B. In fact, in one implementation each of these optical switch modules has 64 optical fibers in and 64 optical fibers out, with these 64 optical fiber pairs fanning out to a great many different line shelves. Different shelves have multiple fiber inputs and outputs. Six parallel optical switch modules 15-1 through 15-6 provide 6 times the data capacity of a single switch module. Other embodiments can have for example, 36 of these modules rather than six.

[0024] Chunks of information are sent individually through optical switch modules 15-1 through 15-N and 16-1 through 16-N and received by IO modules 17 on line shelves at the egress side of router 10. IO module 17 checks the FEC check symbols to validate the accuracy of the data bits within the chunk. It then removes the FEC check symbols and transfers the resulting chunk payloads to packet forwarding module 18-0, 18-1, 18-2, 18-3, or 18-4 as appropriate for each destination address. Similarly, the working one optical switch plane containing optical switch modules 16-1 through 16-N does substantially the same thing in parallel. Thus, working zero and working one optical switch planes perform this process duplicatively and in parallel. This allows the packet forwarding modules on the egress side, such as PFM 18-0, to select those chunk payloads that are error free either from working zero or from working one optical switch plane on a chunk by chunk basis. If there is an error in an optical switch, then egress PFM modules 18-0 through 18-N can identify which working plane, zero or one, is accurate. Consequently errors in a switch are contained and do not ripple out through the network.

[0025] If there are only a few bit errors going through a switch, those errors can be corrected in real time by FEC decoding in IO modules 17. If a path through a

working zero optical switch fails completely, then a path through the working one optical plane can be utilized instead. Further, because each IO module 17 computes the corrupted bits and how many bits were corrected on every path of the system, IO modules 17 provide a detailed fault analysis not only of the failed fiber or optical switch plane, but even down to the level of an individual switch defect, which then can also be isolated. Importantly, the data flowing across for example OS Module 15-1 and the data flowing across OS Module 16-1 in the absence of failures in the system are identical, byte for byte. This provides a hot standby, chunk for chunk.

[0026] After selecting error-free chunk payloads, packet forwarding modules 18-0 through 18-N then re-assemble the chunks into individual IP packets and forward those packets across interface links 104, as previously described.

[0027] In Figs. 1A-1C for the purpose of clarity, corresponding input and output functions are shown on separate circuit cards in separate $\frac{1}{2}$ line shelves 142 and 143 respectively. In some embodiments corresponding input and output functions are combined on a single circuit card in a single line shelf combining $\frac{1}{2}$ line shelves 142 and 143, thereby creating a folded configuration. For example, working input facility module 11-0W and working output facility module 12-0W can be combined on a single physical printed circuit card with two optical connectors, one in and one out. Similarly protect input facility module 11-0P and protect output facility module 12-0P can be combined on a single physical circuit card with two optical connectors, one in and one out. Likewise, input and output packet forwarding modules 13-0 and 18-0 also can be combined on a single physical circuit card in a single line shelf. In a folded configuration, if packet forwarding modules 13-0 and 18-0 share the same physical card, then there is a single card for Channel 0, likewise a single card each for Channels 1, 2, 3, and a fifth card for a Protect channel 4. Because there is a single physical card for input and output functions, then if a card fails, the protection ratio is equal for both input and output modules on that card. In some embodiments internal optics modules 14-0W0 and 17-0W0 similarly share the same physical circuit card, which in the present implementation is contained in the same line shelf 142, 143 with combined input/output facility modules 11, 12 and combined input/output packet forwarding modules 13, 18.

[0028] Fig. 2 is a block diagram illustrating data flow through facility modules 11-0W and 12-0W, for example, in more detail. Facility optical fibers are connected on the left through input and output interfaces 101-0W and 102-0W respectively. In a preferred embodiment shown in Fig. 2, for purposes of illustration input and output facility modules 11-0W and 12-0W occupy the same circuit board in the same line shelf in a folded configuration. In other embodiments, the input and output facility modules 11-0W and 12-0W are located on separate physical circuit cards.

[0029] A signal, e.g., a packet-over-SONET (POS) formatted IP packet, arrives at input 101-OW to a signal processing module 201 typically in a ten-Gbit/sec OC192 SONET datastream. Processing module 201 contains an optical receiver, an optical multiplexer and associated demultiplexer, and a transmitter associated with those. For example, the received signal is demodulated from optical input 101-OW into an electronic signal, and then demultiplexed from a single ten-Gbit-per-second datastream in this example down to a parallel bus at a lower data speed. That parallel bus of signals then leaves module 201 and goes into a processing module 202. Module 202 contains an OC192 demultiplexer, which extracts a single 2.5 Gbit/second OC48 substream out of the OC192 stream and delivers a packet-over-SONET (POS) input to a framer 203-1, which is an industry standard off the shelf component. Likewise, module 202 extracts the other three OC48 substreams and sends these to POS framers 203-2, 203-3, and 203-4 respectively. At this point there are four parallel 2.5 Gbit/sec SONET streams, one to each of four POS framers 203-1 through 203-4, which extract from each OC48 stream the individual IP packets. POS framers 203-1 through 203-4 first have to find the IP packets in the datastream and then have to extract the packets from the SONET continuous datastream. This is done on the four parallel OC48 streams. Once it has removed the packets from the SONET frame, each POS framer 203-1 through 203-4 delivers those packets to a facility ASIC 204-1 through 204-4 respectively.

[0030] The principal function of facility ASICs 204-1 through 204-4 is to send that information to an appropriate packet forwarding module (not shown in Fig. 2), in this case through an interface 103-OW consisting of four parallel interfaces for the four packet streams, or, if directed, to receive packets from an upstream neighboring facility ASIC on an interface 103-4W and switch 103-4W to 103-OW in a protect mode. Otherwise, in a working mode of operation, a facility ASIC sends the information out through interface 103-OW, and information input on 103-4W is directed through cascading protection bus interface 105-OW. The normal sequence is for a facility ASIC to take information from above and switch it below, letting the received traffic pass straight through onto interface 103-OW. All four of facility ASIC switches 204-1 through 204-4 are ganged, such that they operate in parallel. With faster buses, faster framers, or faster facility ASICs, a single ASIC or bus, for example, could perform the above described functions instead of four required at the present state of technology.

[0031] Referring again to Fig. 2, on the egress side facility ASIC 204-1 directs the information packets through output link 211 to Packet-over-SONET framer 203-1, which receives a packet, inserts it into a SONET frame, producing a 2.5 gigabit/second datastream or parallel bus equivalent, and sends that frame to OC192 add/drop multiplexer 202. Multiplexer 202 combines four 2.5 gigabit/second streams from POS framers

203-1 through 203-4, multiplexes them together into a 10 gigabit/second datastream, and delivers them to optical transceiver 201. Transceiver 201 receives the 10 gigabit/second stream, which is formatted as a parallel bus, and multiplexes it into a single datastream, which modulates a laser diode. This produces a SONET ten-gigabit/second optical format, which is transmitted through outbound optical facility interface link 102-OW.

[0032] Fig. 3 is a block diagram illustrating information flow through a typical packet forwarding module 13-0 (18-0). Facility ASICs 301-1 through 301-4 on the ingress side receive packets from facility modules working and protect 11-OW and 11-OP through single links 103-OW0 through 103-OW3. A principal function of facility ASICs 301-1 through 301-4 on the ingress side is to select between the working and the protection facility modules, as represented by the information on, for example, incoming path 103-OW0 or 103-OP0. That selection is made based on the standard SONET criteria for defining if one or both of those incoming facility modules is flawed or failed and also based on any detection of local errors or failures on working facility module 11-OW or protect facility module 11-OP.

[0033] In the egress direction, a principal function of facility ASICs 301-1 through 301-4 is to duplicate the packet stream coming out of egress ASIC 302 and to send that packet stream out across both outgoing paths 104-OW0 and 104-OP0 to facility modules 12-OW and 12-OP (see Fig. 2).

[0034] Packet forwarding engines 306-1 through 306-4 are devices that inspect the packet headers of all of the incoming packets received on any of the selected working or protect facility modules that are associated with this particular packet forwarding module 13-0 (18-0). Based on the inspection of those headers, a determination of the intended destination of each packet can be made. The header information is stored by an ingress ASIC 304 in various queues and lists, which are used to determine for any given packet which output port of the router it should exit, when it should exit, and its relative priority. Actual packet data is stored by ingress ASIC 304 in an external RAM memory 305. Packet forwarding engine 306-1 through 306-4 also determines if any particular packet is intended for a local destination within this particular router and redirects it toward the main control processor of the router instead of transmitting it downstream out one of the output ports of the router to a peer router across the network.

[0035] Ingress ASIC 304, based on the states of the various queues that it maintains and based on the destination addresses of the various packets that are represented by headers in those queues, sends requests through optical transceiver units 308-W and 308-P across optical link 310 (typically multimode ribbon fiber) to the central arbiter (not shown in Fig. 3). The central arbiter determines, based on all of the packets that are being processed through the router in aggregate at any given time, which of the requests from a particular in-

gress ASIC should be granted and when it should be granted for transmission across the optical switch. Grants of those requests return across optical link 310 through transceivers 308-W and 308-P back to ingress ASIC 304. Ingress ASIC 304 uses that grant information to extract packets from memory 305 in the appropriate order and assembles them into chunk payloads. At the appropriate times ingress ASIC 304 sends those chunk payloads across channels 107-00 through 107-03 to internal optics modules 14-0W0 through 14-NW1 (see Fig. 1B).

[0036] On the egress side, information chunk payloads are received from the optical switch matrix indirectly through internal optics modules 17-0W0 through 17-NW1 (see Fig. 1B) across links 108-00 through 108-03 into an egress ASIC 302. Egress ASIC 302 reconfigures the chunks into packets and again stores the packets in a memory 303 in the form of queues and structures. Egress ASIC 302 subsequently reads those packets out again into one of the four facility ASICs 301-1 through 301-4. At the facility ASIC, each of those packet streams is duplicated and sent in tandem to both working and protect facility modules 12-0W and 12-0P.

[0037] A line control processor 307 is primarily responsible for controlling the facility protection switching function by examining the SONET error and failure indications from facility modules 11-0W and 11-0P and also by analyzing the indications that facility ASICs 301-1 through 301-4 develop from those incoming signals. The appropriate switching decisions are made in software and logic and are then implemented by line control processor 307.

[0038] Fig. 4 is a block diagram representing information flow through a typical internal optics module 14 (17), according to an embodiment of the present invention. Internal optics module 14 receives chunk payloads of data via input links 107-00 through 107-04 from packet forwarding modules 13-0 through 13-N (see Fig. 3). An internal optics ASIC 407 selects chunk payloads from those inputs based on grant information that comes back from the central arbiter through each packet forwarding module 13-0 through 13-N. Internal optics ASIC 407 selects which inputs 107-00 through 107-04 will be passed at any point in time to three MUXs 401-1 through 401-3 and out through three 12.5-gigabit-per-second transmitters 403-1 through 403-3 toward the optical switch modules over single mode optical fiber links 109-1 through 109-3. Internal optics ASIC 407 is responsible for encapsulating the chunk payloads with the forward error correcting (FEC) headers and check sums that guarantee that the chunks pass across the optical switch without error, or that if errors occur, they are either corrected or detected. MUXs 401-1 through 401-3 convert input parallel format data to higher bit rate serial data.

[0039] In the egress direction in Fig. 4, optical signals coming in over multimode optical fiber links 110-1 through 110-3 pass through 12.5-gigabit-per-second re-

ceivers 404-1 through 404-3 and into three DEMUXs 402-1 through 402-3. Receivers 404-1 through 404-3 convert the data chunks from optical to electrical bits and DEMUXs 402-1 through 402-3 convert these from a serial bit stream to lower bit rate parallel bit streams. Internal optics ASIC 407 compares the calculated FEC (forward error correction) check sums with the encoded check sums and determines if any errors have occurred across the switch matrix, corrects those errors if possible, and if not, provides alarm and performance monitoring information based on those errors. Internal optics ASIC 407 then strips away the FEC coding from the chunks and passes the resulting chunk payloads from the demux channels out through links 108-00 through 108-04 to packet forwarding modules 18-0 through 18-N.

[0040] In the egress direction, chunk payloads received from internal optics modules 17 are broken down into their original packets by egress ASIC 302 (see Fig. 3). The packets are stored in memory 303 and are then retrieved and delivered at the appropriate time to facility modules 12-0W and 12-0P. Each packet forwarding module 13 packages chunk payloads as described earlier and sends identical streams of chunk payloads to both working 1 and working 0 copies of the optical fabric via internal optics modules (IOMs) 14-0W0 through 14-NW1 (see Fig. 1B). Working 0 copy of the optical switch fabric includes internal optics modules 14-0W0 and 14-1W0, optical switch modules 15-1 through 15-6, and internal optics modules 17-0W0 and 17-1W0, whereas working 1 copy of the optical switch fabric includes internal optics modules 14-0W1 and 14-1W1, optical switch modules 16-1 through 16-6, and internal optics modules 17-0W1 and 17-1W1. For example, IOM 14-0W0 and IOM 14-0W1 each receive simultaneous sequences of chunk payloads from each packet forwarding module 13 that is transmitting through those two IOMs. Similarly, on the egress side each packet forwarding module items 18-0 through 18-N (see Fig. 1C) receives a simultaneous sequence of chunk payloads from IOMs 17-0W0 and 17-0W1, for example. In error-free normal working operation of both optical switch fabrics, the simultaneous sequences of chunk data delivered to each packet forwarding module are identical. In the event of a failure of any kind, either within a chunk or across multiple chunks on either copy zero or copy one of the optical switch fabric, the affected IOM is able to detect that failure based on comparison of the received FEC check sums with the calculated FEC check sums. When a failure on a particular chunk from either working zero or working one copy of the optical switch fabric is detected, the IOM inserts a failure indication downstream toward PFMs 18. This forces PFM 18 to select the error-free chunk data from the alternate copy of the optical switch fabric. This can be done individually for each chunk payload delivered to a particular PFM.

[0041] Referring again to Fig. 3, each packet forwarding module 13 packages chunk payloads as described

earlier and sends identical streams of chunk payloads to both working 1 and working 0 copies of the optical switch fabric via internal optics modules (IOMs) 14-0W0 through 14-NW1 (see Fig. 1B), which encapsulates the chunk payloads into chunks. Working 0 copy of the optical switch fabric (see Fig. 1B) includes internal optics modules 14-0W0 and 14-1W0, optical switch modules 15-1 through 15-6, and internal optics modules 17-0W0 and 17-1W0, whereas working 1 copy of the optical switch fabric includes internal optics modules 14-0W1 and 14-1W1, optical switch modules 16-1 through 16-6, and internal optics modules 17-0W1 and 17-1W1. For example, IOM 14-0W0 and IOM 14-0W1 each receive simultaneous sequences of chunk payloads from each packet forwarding module 13 that is transmitting through those two IOMs. Similarly, on the egress side each packet forwarding module 18-0 through 18-N (see Fig. 1C) receives a simultaneous sequence of chunk payloads from IOMs 17-0W0 and 17-0W1, for example. In error-free normal working operation of both optical switch fabrics, the simultaneous sequences of chunk data delivered to each packet forwarding module are identical. In the event of a failure of any kind, either within a chunk or across multiple chunks on either copy zero or copy one of the optical switch fabric, the affected IOM is able to detect that failure based on comparison of the received FEC check sums with the calculated FEC check sums. When a failure on a particular chunk from either working zero or working one copy of the optical switch fabric is detected, the IOM inserts a failure indication downstream toward PFM 18. This forces PFM 18 to select the error-free chunk data from the alternate copy of the optical switch fabric. This can be done individually for each chunk payload delivered to a particular PFM.

[0042] Referring again to Fig. 4, internal optics ASIC 407 detects any errors or failures of a given chunk on either copy zero or copy one of the switch fabric and inserts appropriate failure indications downstream toward all of the packet forwarding modules connected to it.

[0043] Referring again to Fig. 3, egress ASIC 302 receives those failure indications and selects on a chunk by chunk basis between either the copy zero or the copy one switch fabric. Only error-free chunk payloads from an unfailed switch fabric are inserted into memory and subsequently retrieved and broken out into packets, which are then transmitted toward facility modules 12-0W and 12-0P.

[0044] Fig. 5 is a simplified schematic diagram illustrating a method of fault isolation and diagnostics in a switch fabric, in accordance with embodiments of the present invention. An optical switch fabric 501 in an optical switch module 15 is interconnected through optical links 109-1 through 109-6 with internal optics module 14 (see Fig. 4) on the ingress side and through optical links 110-1 through 110-6 with internal optics module 17 on the egress side. Control and timing signals are received by optical switch module 15 through control and timing

links 503. Also contained in optical switch module 15 are an optical transmitter 502 interconnected with optical switch fabric 501 on the ingress side and an optical receiver 504 interconnected with optical switch fabric 501 on the egress side.

[0045] The fault isolation and diagnostic method is initiated by launching a non-traffic-bearing data chunk from optical transmitter 502 on a predetermined data path through optical switch fabric 501 and then detecting and receiving the data chunk at optical receiver 504. The received data chunk is then examined according to predetermined criteria. If the non-traffic-bearing data chunk satisfies these criteria, then the predetermined data path is error-free. However, if the non-traffic-bearing data chunk fails to satisfy these criteria, then that predetermined data path is faulty. The predetermined criteria typically include forward error correction criteria, although other error checking criteria can be used.

[0046] The method of fault isolation and diagnostics illustrated in Fig. 5 can be applied to multiple duplicated switch fabrics. The non-traffic-bearing data chunk is first encoded with forward error correction or other appropriate coding prior to being launched through the optical switch fabric. The non-traffic-bearing chunk can not only be launched and detected within the optical switch module, but alternatively within other modules, for example internal optics modules. The launch site and receiving site can be on the same or different modules. The method can be applied to non-optical switch fabrics as well.

[0047] Note that while embodiments of the invention have been described in terms of two SONET standards namely OC48 and OC192, alternative implementations of router 10 having an appropriate facility module can operate under other standards.

[0048] Embodiments according to the present invention are designed to protect against all single fault occurrences. Single faults include a single fault of a module, a single fault of a cable, or a single fault of a path. Accordingly, although some double faults are protected against, double faults generally lie beyond the scope of principal objects of the present invention and thus are not in general protected against.

[0049] Although the present invention and its advantages have been described in detail, it should be understood that various changes, substitutions and alterations can be made herein without departing from the spirit and scope of the invention as defined by the appended claims. Moreover, the scope of the present application is not intended to be limited to the particular embodiments of the process, machine, manufacture, composition of matter, means, methods and steps described in the specification. As one of ordinary skill in the art will readily appreciate from the disclosure of the present invention, processes, machines, manufacture, compositions of matter, means, methods, or steps, presently existing or later to be developed that perform substantially the same function or achieve substantially the same result as the corresponding embodiments de-

scribed herein may be utilized according to the present invention. Accordingly, the appended claims are intended to include within their scope such processes, machines, manufacture, compositions of matter, means, methods, or steps.

Claims

1. A communication network comprising:

two duplicated substantially identical switch fabrics, including a first switch fabric comprising a first $N \times M$ crossbar switch and a second switch fabric comprising a substantially identical second $N \times M$ crossbar switch, wherein N is the number of ingress ports and M is the number of egress ports of each respective $N \times M$ crossbar switch;

wherein said first and second crossbar switches are connected in substantially identical parallel data paths, such that each ingress port of said first and said second crossbar switch is interconnected with a data launching module, and each egress port of said first and said second crossbar switch is interconnected with a data receiving module; and

wherein the geometry of said switch fabrics is folded, such that a data launching module and a data receiving module occupy the same physical circuit card.

2. The communication network of claim 1 wherein said crossbar switches are optical switches.

3. The communication network of claim 2 wherein said optical switches are interconnected with said data launching and said data receiving modules through optical fibers.

4. The communication network of claim 1 wherein said data launching module and said data receiving module are internal optics modules.

5. The communication network of claim 1 wherein said data launching module is interconnected with an ingress data forwarding module, and said data receiving module is interconnected with an egress data forwarding module.

6. The communication network of claim 1 further comprising a router system, said router system incorporating said first and said second switch fabrics.

7. A communication network comprising:

two duplicated substantially identical switch

fabrics, including a first switch fabric comprising a first $N \times M$ optical crossbar switch and a second switch fabric comprising a substantially identical second $N \times M$ optical crossbar switch, wherein N is the number of ingress ports and M is the number of egress ports of each respective $N \times M$ crossbar switch;

wherein said first and second optical crossbar switches are connected in substantially identical parallel data paths, such that each ingress port of said first and said second optical crossbar switch is interconnected with a data launching module, and each egress port of said first and said second optical crossbar switch is interconnected with a data receiving module; and

wherein each of said first and said second optical crossbar switches is configured to pass information at a data rate of approximately 12.5 gigabits per second.

8. The communication network of claim 7 wherein the geometry of said switch fabrics is folded, such that a data launching module and a data receiving module occupy the same physical circuit card.

9. The communication network of claim 7 wherein said optical crossbar switches are interconnected with said data launching modules and said data receiving modules through optical fibers.

10. The communication network of claim 7 wherein said data launching module and said data receiving module are internal optics modules.

11. The communication network of claim 7 wherein said data launching module is interconnected with an ingress data forwarding module, and said data receiving module is interconnected with an egress data forwarding module.

12. The communication network of claim 7 further comprising a router system, said router system incorporating said first and said second switch fabrics.

13. A method of switch fabric protection comprising:

simultaneously launching parallel duplicate data streams through two duplicated substantially identical switch fabrics, including launching a first data stream through a first switch fabric comprising a first $N \times M$ crossbar switch, and launching a substantially identical second data stream through a second switch fabric comprising a substantially identical second $N \times M$ crossbar switch, wherein N is the number of ingress ports and M is the number of egress ports of each respective $N \times M$ crossbar switch;

receiving said parallel duplicate data streams after passing simultaneously through said first and said second switch fabric;

examining said received duplicate data streams in accordance with predetermined selection criteria;

if either of said duplicate data streams satisfies said criteria and the other said duplicate data stream does not satisfy said criteria, then selecting said duplicate data stream that satisfies said criteria and discarding said duplicate data stream that does not satisfy said criteria; and

if both of said duplicate data streams satisfy said criteria, then arbitrarily selecting one of said duplicate data streams and arbitrarily discarding the non-selected duplicate data stream.

14. The method of claim 13 wherein said duplicate data streams comprise duplicate sequences of data structures, wherein said data structure is selected from the group consisting of data packets and substantially fixed size data chunks.

15. The method of claim 14 wherein said data structures are encapsulated before said launching with a code selected from the group consisting of forward error correction code and cyclic redundancy code.

16. The method of claim 15 wherein said selecting of said data stream is performed on a structure-by-structure basis in accordance with said code encapsulated with said data structure.

17. The method of claim 13 wherein data delivery by said data streams is not interrupted by an occurrence selected from the group consisting of malfunction, failure, removal, and replacement of one of said two duplicated substantially identical switch fabrics.

18. The method of claim 13 wherein said examining is performed at an egress internal optics module interconnected with an egress port of each of said duplicated substantially identical switch fabrics

19. The method of claim 18 wherein said selecting and said discarding are performed at an egress packet forwarding module interconnected with said egress internal optics module.

20. The method of claim 13 wherein said first and said second switch fabric are incorporated into a router system.

21. The method of claim 18 wherein said egress internal optics module is interconnected with said egress port through an optical fiber.

22. The method of claim 16 wherein said forward error correction corrects errors in said data structures.

23. The method of claim 16 wherein said forward error correction detects uncorrectable errors in said data structures.

24. A method of fault isolation and diagnostics in a switch fabric comprising:

launching a non-traffic-bearing data structure on a predetermined data path through said switch fabric from a first module interconnected with said switch fabric;

detecting and receiving said non-traffic-bearing data structure at a predetermined second module interconnected with said switch fabric;

examining said received non-traffic-bearing data structure in accordance with predetermined criteria;

if said non-traffic-bearing data structure satisfies said criteria, then determining that said predetermined data path is error-free; and

if said non-traffic-bearing data structure fails to satisfy said criteria, then determining that said predetermined data path is faulty.

25. The method of claim 24 wherein said switch fabric comprises multiple duplicated switch fabrics.

26. The method of claim 24 wherein said predetermined criteria comprise forward error correction.

27. The method of claim 24 wherein said switch fabric comprises an optical crossbar switch.

28. The method of claim 24 wherein said predetermined data path comprises an optical fiber cable.

29. The method of claim 24 wherein said non-traffic-bearing data structure is a substantially fixed size diagnostic chunk.

30. The method of claim 29 wherein identifications representing physical identities of said respective first module and of said second module are encoded into said diagnostic chunk prior to said launching; and said predetermined criteria comprise coincidence between the physical identities of said respective first and second modules and said respec-

tive encoded identifications of said first and second modules.

31. The method of claim 24 wherein said switch fabric, said first module, and said second module are incorporated within a router system. 5

10

15

20

25

30

35

40

45

50

55






FIG. 1A





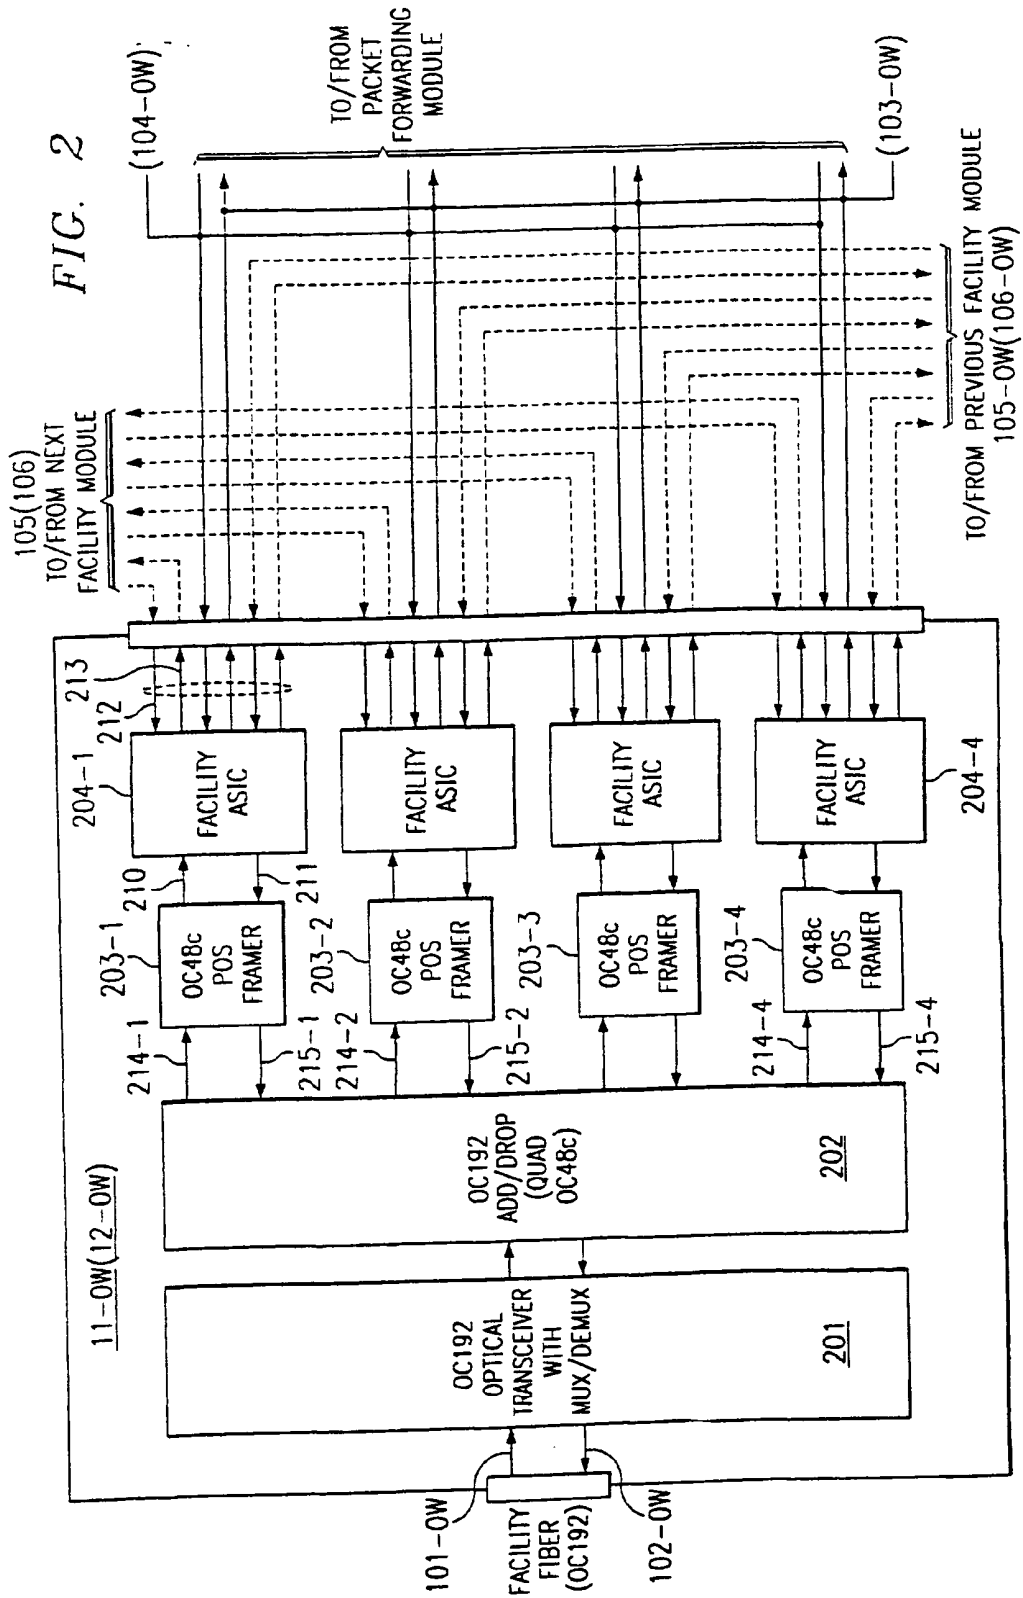



FIG. 3








FIG. 4

FIG. 5



THIS PAGE BLANK (USPTO)



Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 1 202 597 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
30.06.2004 Bulletin 2004/27

(51) Int Cl.7: H04Q 11/00, H04L 12/56

(43) Date of publication A2:
02.05.2002 Bulletin 2002/18

(21) Application number: 01309176.4

(22) Date of filing: 30.10.2001

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE TR
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: 31.10.2000 US 703064

(71) Applicant: Chiaro Networks Ltd.
Richardson, Texas 75080 (US)

(72) Inventors:
• McDermott, Thomas C.III
Plano, Texas 75094 (US)
• Blackmon, Harry C.
Plano, Texas 75023 (US)
• Brewer, Tony M.
Plano Texas 75093 (US)


- Dozier, Harold W.
Dallas, Texas 75248 (US)
- Kleiner, Jim
Dallas Texas 75240 (US)
- Palmer, Gregory S.
Plano, Texas 75025 (US)
- Shaw, Keith W.
Plano, Texas 75023 (US)
- Traylor, David
Richardson, Texas 75080 (US)
- Walker, Dean E.
Plano, Texas 75025 (US)



(74) Representative: Jackson, Richard Eric et al
Carpmaels & Ransford,
43 Bloomsbury Square
London WC1A 2RA (GB)

(54) Router-Switch fabric protection using forward error correction

(57) Instead of alternatively utilizing only one fabric or the other fabric of a redundant pair, both fabrics simultaneously transmit duplicate information, such that each packet forwarding module (PFM) receives the output of both fabrics simultaneously. In real time, an internal optics module (IOM) analyzes each information chunk coming out of a working zero switch fabric; simultaneously examines a parallel output of a working one duplicate switch fabric; and compares on a chunk-by-chunk basis the validity of each and every chunk from both switch fabrics. The IOM does this by examining forward error correction (FEC) check symbols encapsulated into each chunk. FEC check symbols allow correcting a predetermined number of bit errors within a chunk. If the chunk cannot be corrected, then the IOM provides indication to all PFMs downstream that the chunk is defective. Under such conditions, the PFMs select a chunk from the non-defective switch fabric. Under error-free normal conditions, however, the PFMs select a chunk arbitrarily from a default switch fabric. In this way, each chunk in real time is selected from a non-defective source and is thus guaranteed to be error free. Accordingly, if a switch fabric fails, no information chunks are

lost anywhere in the system.







European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 01 30 9176

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
X	WO 94/27391 A (ITALTEL SPA ; COLLIVIGNARELLI MARCO (IT)) 24 November 1994 (1994-11-24)	1,5,6,8, 11-17, 20,22,23	H04Q11/00H04L12/ H04Q11/00
Y	* page 2, line 21 - page 4, line 2 *	2-4,7,9, 10,18, 19,21	
	* page 5, line 33 - page 8, line 25 *		
	* page 9, line 19 - page 16, line 30 *		
Y	BOSTICA B ET AL: "THE ROLE OF ELECTRONICS AND PHOTONICS IN THE ARCHITECTURAL EVOLUTION OF SWITCHING SYSTEMS" ISS '97. WORLD TELECOMMUNICATIONS CONGRESS. (INTERNATIONAL SWITCHING SYMPOSIUM). GLOBAL NETWORK EVOLUTION: CONVERGENCE OR COLLISION? TORONTO, SEPT. 21 - 26, 1997, ISS. WORLD TELECOMMUNICATIONS CONGRESS. (INTERNATIONAL SWITCHING SYMPOSIUM), TORONTO, P, vol. 2, 21 September 1997 (1997-09-21), pages 565-571, XP000704512	2-4,7,9, 10,18, 19,21	
A	* page 565, left-hand column, line 44 - right-hand column, line 22 *	1,5,6,8, 11-17, 20,22,23	H04Q H04L
	* page 566, left-hand column, line 9 - line 31 *		
	* page 567, right-hand column, line 35 - page 570, left-hand column, line 1 *		
	* figure 8 *		
----- -/--			
The present search report has been drawn up for all claims			
Place of search Munich		Date of completion of the search 16 February 2004	Examiner AVILES MARTINEZ, M
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons</p> <p>& : member of the same patent family, corresponding document</p>			

EPO FORM 1503 03.02 (P04C01)



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 01 30 9176

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.7)
A	<p>AWEYA J: "On the design of IP routers Part 1: Router architectures" JOURNAL OF SYSTEMS ARCHITECTURE, ELSEVIER SCIENCE PUBLISHERS BV., AMSTERDAM, NL, vol. 46, no. 6, April 2000 (2000-04), pages 483-511, XP004190486 ISSN: 1383-7621</p> <p>* page 493, paragraph 3.3; figure 6 *</p> <p>-----</p>	1-23	
A	<p>EP 1 037 492 A (FURUKAWA ELECTRIC CO LTD) 20 September 2000 (2000-09-20)</p> <p>* page 4, line 13 - line 54 *</p> <p>* page 5, line 49 - line 56 *</p> <p>* page 6, line 12 - page 8, line 54 *</p> <p>-----</p>	1-23	
<p>The present search report has been drawn up for all claims</p>			<p>TECHNICAL FIELDS SEARCHED (Int.Cl.7)</p>
Place of search		Date of completion of the search	Examiner
Munich		16 February 2004	AVILES MARTINEZ, M
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone</p> <p>Y : particularly relevant if combined with another document of the same category</p> <p>A : technological background</p> <p>O : non-written disclosure</p> <p>P : intermediate document</p>		<p>T : theory or principle underlying the invention</p> <p>E : earlier patent document, but published on, or after the filing date</p> <p>D : document cited in the application</p> <p>L : document cited for other reasons</p> <p>& : member of the same patent family, corresponding document</p>	

EPO FORM 1503 03/02 (P4/G01)



European Patent
Office

Application Number

EP 01 30 9176

CLAIMS INCURRING FEES

The present European patent application comprised at the time of filing more than ten claims.

- ☐ Only part of the claims have been paid within the prescribed time limit. The present European search report has been drawn up for the first ten claims and for those claims for which claims fees have been paid, namely claim(s):
- ☐ No claims fees have been paid within the prescribed time limit. The present European search report has been drawn up for the first ten claims.

LACK OF UNITY OF INVENTION

The Search Division considers that the present European patent application does not comply with the requirements of unity of invention and relates to several inventions or groups of inventions, namely:

see sheet B

- ☐ All further search fees have been paid within the fixed time limit. The present European search report has been drawn up for all claims.
- ☐ As all searchable claims could be searched without effort justifying an additional fee, the Search Division did not invite payment of any additional fee.
- ☐ Only part of the further search fees have been paid within the fixed time limit. The present European search report has been drawn up for those parts of the European patent application which relate to the inventions in respect of which search fees have been paid, namely claims:
- ☒ None of the further search fees have been paid within the fixed time limit. The present European search report has been drawn up for those parts of the European patent application which relate to the invention first mentioned in the claims, namely claims:

1-23



European Patent
Office

LACK OF UNITY OF INVENTION
SHEET B

Application Number
EP 01 30 9176

The Search Division considers that the present European patent application does not comply with the requirements of unity of invention and relates to several inventions or groups of inventions, namely:

1. claims: 1-23

Claims 1 to 23 refer to a communications network and a method of operating it using duplicated switch fabrics.

2. claims: 24-31

Claims 24 to 31 refer to a method of fault isolation and diagnostics in a switch fabric.

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 01 30 9176..

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

16-02-2004

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
WO 9427391	A	24-11-1994	IT 1272471 B	23-06-1997
			WO 9427391 A1	24-11-1994
			EP 0699365 A1	06-03-1996

EP 1037492	A	20-09-2000	JP 2000269890 A	29-09-2000
			JP 2001021929 A	26-01-2001
			EP 1037492 A2	20-09-2000
			US 6477288 B1	05-11-2002
			CA 2312499 A1	12-01-2001

EPO FORM P0459

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82

THIS PAGE BLANK (USPTO)

**ORIGINAL
NO MARGINALIA**